

Robust Pedestrian Tracking in Crowd Scenarios Using an Adaptive GMM-based Framework

Shuyang Zhang¹, Di Wang², Fulong Ma³, Chao Qin¹, Zhengyong Chen¹ and Ming Liu³

Abstract—In this paper, we address the issue of pedestrian tracking in crowd scenarios. People in close social relationships tend to act as a group which is a great challenge to individually discriminate and track pedestrians on a LiDAR system. In this paper, we integrally model groups of people and track them in a recursive framework based on Gaussian Mixture Model (GMM). The model is optimized by an extended Expectation-Maximization (EM) algorithm which can adaptively vary the number of mixture components over scans. Experimental results both qualitatively and quantitatively indicate the reliability and accuracy of our tracker in populated scenarios.

I. INTRODUCTION

A. Motivation

Multi-object tracking (MOT) is crucial for mobile robotics applications. The host robot monitors the complicated motion of its surroundings and provides reliable motion estimates for subsequent decision-making module. It is especially intractable in crowd scenarios, such as factory and campus, where pedestrians are the major participants and have special interactive behaviors [1] [2] [3]. People in intimate social relationships tend to be in close physical position and normally crowd or walk as a group in public.

The first challenge of tracking groups of people is to extract individuals. As shown in Fig. 1, friends walk as a group and have small physical intervals. It is a hardship for spatial segmentation methods to cluster each individual from a LiDAR scan. The second challenge is to keep association consistency under the interaction of other group members. The Euclidean distance between two close pedestrians keeps changing over scans. Sometimes, they can be separately segmented but in some cases are wrongly considered as one. This uncertainty in segmentation makes data association ambiguous and ultimately causes the degradation of tracking performance.

Tracking groups of people in LiDAR perception is researched for decades [2] [3] [4]. An inspiring thought is to model individuals and people groups separately [2] [4] [5] and they proposed that modeling a group of people as a single object is more efficient. Pedestrians and people groups are respectively clustered and modeled in reasonable representation, such as contours [2] and Gaussian shapes

¹ Shuyang Zhang, Chao Qin and Zhengyong Chen are with Unity-Drive technology Inc, Shenzhen, China {shuyang.zhang, chao.qin, chenzhengyong}@unity-drive.com.

² Di Wang is with Institute of Artificial Intelligence and Robotics, Xian Jiaotong University, China {de2wang}@stu.xjtu.edu.cn.

³ Fulong Ma and Ming Liu are with the Robotics and Multi-Perception Laboratory, Robotics Institute, The Hong Kong University of Science and Technology, Hong Kong SAR, China. {fmaaf, eelium}@ust.hk.



Fig. 1. Hybrid view of a crowd scenario. The point cloud is collected by four RS-LiDAR-16 scanners on our intelligent delivery robot. People marked by the yellow box are in close proximity and it is challenging to cluster each individual using spatial segmentation methods.

[5]. Afterward, they are tracked within Bayesian filtering frameworks. Lau *et al.* [2] also creatively discussed the interconversion and interaction between individuals and groups.

Similar to the work of Mucientes *et al.* [5], we model pedestrians by using Gaussian shapes. The difference is that we also model the entire LiDAR scan as a Gaussian Mixture Model (GMM) combined with a uniform clutter model. The number of Gaussian components, which represents the number of pedestrians in the scene, is supported by the prior model and it can be adaptively updated using our extended Expectation-Maximization (EM) algorithm.

B. Contribution

In this paper, we addressed this issue of individual pedestrian tracking in crowd scenarios. The contributions of this paper are summarized as follows:

- A GMM-based framework focused on tracking individuals in crowd scenarios.
- An extended EM algorithm for GMM parameter learning which adaptively changes the number of Gaussian components.
- Quantitative and qualitative evaluation with real scene data which shows robust tracking performance under people's interaction.

C. Organization

The remainder of this paper is structured as follows. Sect. II systematically introduces state of the art on LiDAR-based tracking. We give an overview of our system in Sect. III. The mixture model, extended EM algorithm and detailed tracking pipeline are respectively presented in Sect. IV before we evaluate the tracker in Sect. V. Lastly, the conclusion and future work are summarized in Sect. VI.

II. RELATED WORKS

Multi-object tracking (MOT) for groups of people are extensively studied in camera-based [6] [7] [8] [9], LiDAR-based [2] [3] [4] and fusion-based [10] [11] system. In this paper, we only focus on the LiDAR-based MOT and categorize them as model-free and model-based methods.

A. Model-free Methods

Model-free methods can generally adapt to objects with various shapes and sizes. Firstly, they extract objects from a LiDAR scan in the form of clusters using spatial cues, such as distance [12] [13] and angle [14] or using motion cues [15] [16]. Later, observation is aligned with existing tracks via their position or feature similarity. Finally, the track state of each object is independently or jointly updated under a Bayesian filtering framework. These methods using spatial cues are implemented under an intuitive pipeline. However, the performance of the data association is highly related to the prior segmentation procedures. Under- and over-segmentation frequently occur and alignment ambiguity is brought into data association, which ultimately degrades tracking performance.

Methods using motion cues [15] distinguish objects via their movement difference and is rarely affected by association ambiguity. However, pedestrians in a group usually share similar movement and motion-based methods can not perfectly separate individuals.

B. Model-based Methods

In model-based methods, objects are detected and associated with the knowledge of prior geometric models. For vehicle tracking, Petrovskaya *et al.* [17] used a rectangle measurement model in their likelihood-field-based framework while Fortin *et al.* [18] proposed a polyhedron-like one. For pedestrians tracking, Shackleton *et al.* [3] proposed a pedestrian surface model. They collect a set of contour descriptions on people’s surface over scans and align them using a surface matching technique. Spinello *et al.* [19] proposed a sliced pedestrian model. They subdivide a person into multiple layers defined by height and learn a specialized classifier for each layer. Mucientes *et al.* [5] modeled groups of people as Gaussian shapes. They consider data association and clusters split in a Multiple Hypothesis Tracking (MHT) framework.

Our work is a model-based one and we assume that people can be represented as Gaussian shapes. The difference is that we also integrally model an entire LiDAR scan by a mixture model. The number of mixture components is adaptively updated over scans.

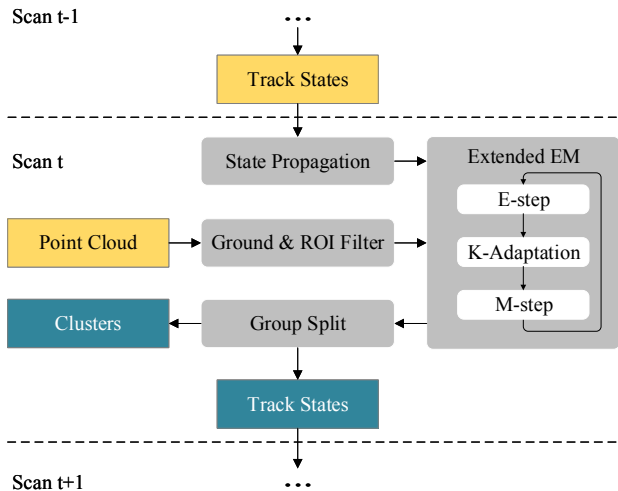


Fig. 2. System pipeline of the proposed method.

III. OVERVIEW

The pipeline of our pedestrian tracking framework is shown in Fig. 2. For each scan, the system inputs are the previous track states and the incoming point cloud. After the state propagation and point cloud filtering, the inputs are used to initialize a GMM-based mixture model which is then optimized by our extended EM algorithm. Group split is afterward considered and the final outputs are the cluster results and the updated track states.

We assume that our robot benefits from a good localization system, so the uncertainty of ego-motion can be negligible over the time scale of the whole track. We use RS-LiDAR-16 scanners, which have poor vertical resolution of 1° and are not sufficient for 3D segmentation. Meanwhile, our intelligent delivery robot is deployed on the road surface and objects in the workspace can be constrained on the 2D plane. Like other autonomous driving applications [20] [21], points are converted to 2D representation in the top-down coordinate system.

IV. METHODOLOGY

A. GMM-Based Scan Modeling

After ground and ROI filtering, a LiDAR scan can be represented as

$$\mathcal{X} = [\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_N] \in \mathbb{R}^{2 \times N}, \quad (1)$$

where N denotes the number of points in 2D format and N changes over scans. Inspired by Horaud *et al.* [22], we model \mathcal{X} as a Gaussian mixture model with a uniform distribution

$$p(\mathbf{x}_i) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}_i | \boldsymbol{\mu}_k, \boldsymbol{\sigma}_k) + \frac{\pi_0}{h}. \quad (2)$$

The number of the Gaussian components K represents the amount of observed objects in the scan. $\boldsymbol{\mu}_k \in \mathbb{R}^2$ and $\boldsymbol{\sigma}_k \in \mathbb{R}^{2 \times 2}$ are the parameters of the k -th Gaussian component and

π_k denotes the mixture coefficient where $\sum_{k=0}^K \pi_k = 1$. The parameters of our model can be denoted as

$$\Theta = \{ \{ \pi_k \}_{k=0}^K, \{ \mu_k \}_{k=1}^K, \{ \sigma_k \}_{k=1}^K \}. \quad (3)$$

The clutter is modeled as a uniform distribution and its coefficient π_0 is manually set to control the proportion of outliers. h represents the area of 2D workspace which contains all the data points. π_0 and h jointly determine the stability of the Gaussian shapes and we will give a detailed discussion in Sect. V-B.

Traditional methods choose K empirically. It is difficult to directly determine the number of pedestrians in a scan. We initialize the first scan by DBSCAN [23], a spatial segmentation method which does not need a selected K . The GMM parameters in a scan are propagated and used for the initialization of the mixture model in the next scan.

B. Extended Expectation-Maximization algorithm

The parameter learning of the Gaussian mixture model is solved via an extended EM algorithm which can be subdivided into 3 parts: Expectation, K-Adaptation and Maximization. In each iteration, the Expectation step updates the posteriors of the hidden variables. The K-Adaptation step changes the number of the Gaussian components K in the model. The Maximization step optimizes the model parameters Θ .

The hidden variables are defined as $\mathbf{Z} = \{z_i\}_{i=1}^N$, where $z_i \in \{0, 1, 2, \dots, K\}$ indicates the component that \mathbf{x}_i belongs to, and \mathbf{Z} can also be considered as the segmentation results of the point cloud \mathcal{X} . The posterior $\gamma_{ik} = p(z_i = k | \mathbf{x}_i)$ denotes the probability that point \mathbf{x}_i aligns to the k -th component. The set of γ_{ik} is an $N \times (K + 1)$ matrix denoted as Γ .

1) *Expectation*: The posteriors Γ are updated by the mixture model parameters Θ . For each posterior

$$\gamma_{ik} = \begin{cases} \frac{1}{\eta} \pi_k \mathcal{N}(\mathbf{x}_i | \mu_k, \sigma_k), & k \neq 0, \\ 1 - \sum_{j=1}^K \frac{1}{\eta} \pi_j \mathcal{N}(\mathbf{x}_i | \mu_j, \sigma_j), & k = 0, \end{cases} \quad (4)$$

where $\eta = \sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{x}_i | \mu_j, \sigma_j) + \frac{\pi_0}{h}$.

2) *K-Adaptation*: In this step, we consider the variation of the component number K . We prune the Gaussian components of exiting objects and append the incoming ones.

For pruning the exiting objects, we use the posteriors Γ updated in the Expectation step. A component needs to be pruned when its point size is smaller than a threshold n_{min} . The number of points that belong to the k -th Gaussian component is defined as

$$n_k = \sum_{i=1}^N \mathbb{I}(\arg \max_j (\gamma_{ij}) = k). \quad (5)$$

$\mathbb{I}(\cdot)$ is an indicator function that it equals to 1 if the input is true and 0 otherwise. The number of pruned components is

$$K_{prune} = \sum_{k=1}^K \mathbb{I}(n_k < n_{min}). \quad (6)$$

For appending the incoming objects, we select the points marked as outliers

$$\mathcal{X}_0 = \left\{ \mathbf{x}_i | \arg \max_j \gamma_{ij} = 0, i = 1, \dots, N \right\}. \quad (7)$$

We segment the outliers \mathcal{X}_0 by DBSCAN and the K_{append} new clusters are initialized as new Gaussian components.

After pruning and appending, the number of components in the mixture model is

$$K_{new} = K - K_{prune} + K_{append}. \quad (8)$$

The posteriors Γ turns to a $N \times (K_{new} + 1)$ matrix. The columns of the trimmed components are deleted and those of the new ones are initialized as $1/K_{new}$. The coefficients $\{ \pi_k \}_{k=0}^K$ are updated by Γ and then re-normalized.

3) *Maximization*: The mixture model parameters Θ are separately updated by the posteriors Γ . The log-likelihood function of Eq.2 is

$$\mathcal{L}(\mathcal{X} | \Theta) = \sum_{i=1}^N \ln \left(\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}_i | \mu_k, \sigma_k) + \frac{\pi_0}{h} \right). \quad (9)$$

The GMM parameters are respectively calculated by standard optimization techniques [24] and we do not restate the details. One point of which needs to be reminded is the update of the mixture coefficients. As a uniform distribution is embedded into the mixture model, the sum of the Gaussian coefficients needs to be subjected to $1 - \pi_0$.

C. Tracking Cycle

This part describes the steps in a cycle of our pedestrian tracker. The details are shown in Fig. 2. We use a tracking-before-detection technique. Firstly, our tracker propagates people's position to the current scan. Afterward, it models the scan as a GMM and performs segmentation by using our extended EM algorithm. Finally, split and merge of group people are considered.

1) *Track State*: We do not implement a general probabilistic filter for pedestrian tracking. The principal reason is that people's motion can not be simply considered as physics-based models [1]. Interactive behaviors in a populated environment make people's movement unpredictable.

Our tracker uses the Gaussian shapes in Sect. IV-A directly. We define a track \mathcal{T} as a sequence of states derive from the same object over scans where $\mathcal{T}_n = \{ \mathcal{S}_{n,t} \}_{t=1}^T$. For each state in a track, $\mathcal{S}_n = \{ \mu_n, \sigma_n, \mathbf{v}_n \}$. $\mu_n = (x, y)^\top$ is the position vector and $\sigma_n \in \mathbb{R}^{2 \times 2}$ represents the shape of the target pedestrian. The velocity part $\mathbf{v}_n = (v_x, v_y)^\top$ is differentially calculated and smoothed by a Kalman smoother. We need to clarify that in our tracking pipeline, the subscripts lose their particular meaning in Sect. IV-A and are only used to indicate specified track.

2) *Propagation*: At the beginning of a tracking cycle, we propagate previous track states to the current scan using constant velocity motion model and make them the initial

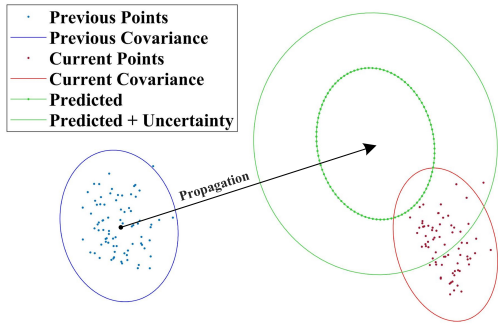


Fig. 3. State propagation for an object. After adding uncertainty to the propagation step, points are more likely to converge to close-by Gaussian components instead of the clutter. The optimization of GMM may converge incorrectly when directly using the covariances from the former scan. The ellipse denotes the 1-sigma Gaussian boundary.

values for our GMM-based modeling. The μ_n and σ_n in track state is predicted as

$$\begin{cases} \bar{\mu}_n = \mu_n + \mathbf{v}_n \delta t, \\ \bar{\sigma}_n = \sigma_n + \mathbf{R} \delta t, \end{cases} \quad (10)$$


where δt denotes the interval between adjacent scans. $\mathbf{v}_n \in \mathbb{R}^2$ is the estimated velocity of the n -th tracked pedestrian while $\mathbf{R} \in \mathbb{R}^{2 \times 2}$ represents a pedestrian's maximum motion uncertainty in a unit time.

As shown in Fig. 3, we add this motion uncertainty directly into a pedestrian's shape model to improve the stability of EM initialization. In Expectation step, γ_{ik} becomes larger accompanied with the increase of σ_k and surrounding points are more likely related to the Gaussian components than outliers. \mathbf{R} takes effect on σ_k only once for a scan. After the first Maximization step, the covariance will converge to the proper size.

3) *Joint Segmentation and Data Association*: We use a tracking-before-detection technique to perform segmentation and data association simultaneously. Segmentation results are derived by the optimized mixture model while data association is realized during the propagation and the update of model parameters over scans. The details are described in Sect. IV-B.

4) *Tracking Management*: The entrance of new objects and the exit of existing ones are also considered by our extended EM algorithm in Sect. IV-B.

We also use some traditional tracking techniques in our tracker. When our tracker loses an object, we do not discard



(a) LiDARs configuration on our platform. (b) A glance of *hard* dataset with ground truth annotation.

Fig. 4. Our intelligent delivery robot and pedestrian tracking dataset.

it immediately but continue a prediction for 5 scans. Gating is also realized in our tracker. It is adjusted by the clutter model and the details will be specifically discussed in Sect. V-B.

D. Group Merge and Split

In this part, we discuss the merge of individuals and the split of people groups.

1) *Merge*: When two pedestrians reunion, they can be separated no matter how close they are because they are already modeled as two Gaussian components in previous scans. These two closed-by persons may not be perfectly divided in pointwise level, but they are able to ensure tracking consistency.

2) *Split*: When two persons are initially modeled as one and walk away from each other in subsequent scans, they are still modeled as one Gaussian distribution, even though they can be spatially separated.

We implement an additional DBSCAN algorithm for the whole scan and we use the segmentation result to validate our GMM-based segmentation. The segmentation output from DBSCAN is represented as $\mathbf{C} = \{c_i\}_{i=1}^N$, which has a similar representation as \mathbf{Z} . We define the point index of the j -th component in our model as $\mathbf{I}_{\mathbf{Z}_j} = \{i | z_i = j\}$. For each cluster m in \mathbf{C} , we calculate their ratio

$$r = \frac{|\mathbf{I}_{\mathbf{C}_m} \cap \mathbf{I}_{\mathbf{Z}_j}|}{|\mathbf{I}_{\mathbf{Z}_j}|}, \quad (11)$$

where $|\cdot|$ indicates the cardinality of the index set. If $r_{min} < r < 1$ and $|\mathbf{I}_{\mathbf{C}_m} \cap \mathbf{I}_{\mathbf{Z}_j}| > n_{min}$, we model each $\mathbf{I}_{\mathbf{C}_m} \cap \mathbf{I}_{\mathbf{Z}_j}$ as a new Gaussian component.

E. Initialization

An incoming scan is pre-processed by a ground and ROI filter. Ground points are filtered by a Gaussian-Process-Regression-based method [25]. The 3D ROI filter removes points out of our robot's motion space. For the initialization of our tracker's first scan, DBSCAN is implemented to perform a spatial segmentation. DBSCAN is also used in our extended EM algorithm and group people split module.

V. EXPERIMENTS

In this section, we introduce the datasets and the details of baseline methods. We also detail the implementation parameters and evaluate our method in tracking performance, cardinality estimates and timing.

A. Datasets

For experimental evaluation, we collected two datasets with different challenge levels. The data is captured by our intelligent delivery robot (Fig. 4(a)) in a populated area of an industrial park. Our LiDAR system is composed of four spinning RS-LiDAR-16 scanners. All LiDARs are hardware synchronized and calibrated [26]. The details of these two datasets are as follows:

- *easy*: A specific dataset that contains 6 independent sequences. There are only several tracks per sequence,

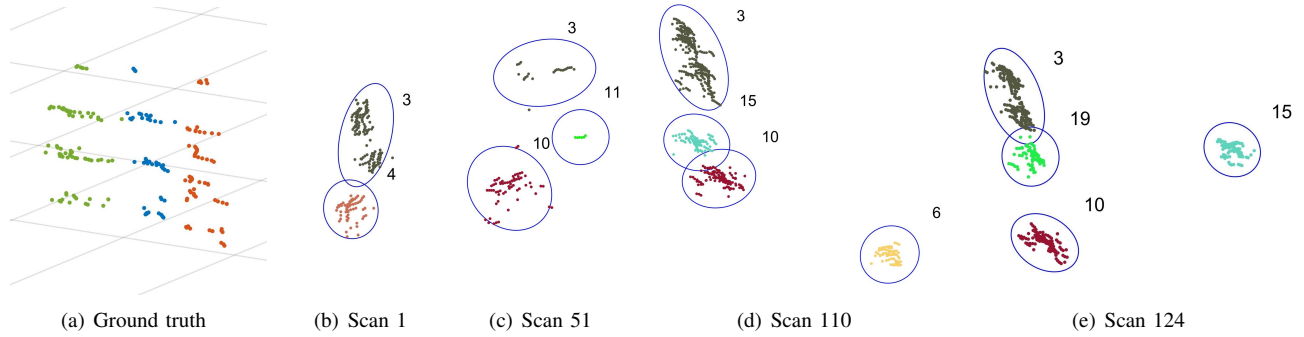


Fig. 5. Qualitative analysis of our tracker. (a) The segmentation ground truth of three pedestrians in the first scan. (b) Our tracker is initialized by DBSCAN which clusters two individuals as one object (No.3). (c) Occlusion makes a wrong split and creates a new observation for clutter (No.11). It vanishes in 5 scans but causes an ID switch from No.4 to No.10. (d) A person (No.15) passes across the people of three. Our tracker handles this interactive condition effectively. (e) Another one (No.19) walks by the group. They are too close that their points mix together but our tracker can still separate them. It is important to point out that our tracker can not subdivide the two people on track No.3. Their distance never exceeds the split threshold. However, it is reasonable to track them as one because they share the same motion.



Fig. 6. The Gaussian 1-sigma ellipse with a different choice of π_0 . The man is waving his hand and the shape of point cloud has a great change over scans. When using a larger π_0 , points away from the Gaussian center are snatched by the clutter model.

but all of these tracks last more than 100 scans. Group merge and object entrance happen but in a limited number.

- *hard*: A general sequence on a populated street, which lasts for 395 scans and contains 36 tracks and 2667 pedestrian observations. The dataset is more complicated where group merge and split, occlusion, object entrance and exit are fully contained.

We use our initialization procedure in Sect. IV-E to filter ground points and the ones out of ROI. The remained points are hand-labeled for each individuals.

B. Implementation Details

The values of π_0 and h are crucial for modeling. As shown in Fig. 6, the component π_0/h in our model represents the signal strength of clutter and it affects the actual range of each Gaussian model in segmentation. It can be considered as a gating function which contributes to the association between model components and raw points. We assume that the effective range of RS-LiDAR-16 for detecting pedestrians is $\pm 10m$, so the area of our workspace in the clutter model $h = 400m^2$. Combined with the selection of h , we set $\pi_0 = 0.001$.

The rest of the parameters are determined empirically. EM optimization terminates with a maximum iteration $n_{max} = 10$ and a convergence criterion that the average variation of μ_k is under $0.01m$. In K-Adaptation step, $n_{min} = 3$. In DBSCAN algorithm, the minimum number of points is the same as n_{min} and the radius threshold in Sect. IV-E and Sect. IV-B is $0.2m$.

The radius threshold of DBSCAN is set as $0.5m$ in Sect. IV-D to balance under-segmentation with occlusions cases. In tracking cycle, pedestrians motion uncertainty $\mathbf{R} \in \mathbb{R}^{2 \times 2}$ and we set $\mathbf{R} = \text{diag}([1.0, 1.0])$. In group split, the ratio threshold for split people group $r_{min} = 0.3$.

C. Baseline Methods

We implement a segmentation method of Wang *et al.* [12] which is based on Euclidean Minimum Spanning Tree (EMST). As the background points in our dataset have been filtered, we do not implement foreground extraction in their original work. The threshold for breaking the EMST is set as $0.2m$. We make an additional data association module achieved by the Munkres algorithm [27] and each cluster is associated via the similarity of their gravity center.

Spectral Clustering (SC) [28] uses high dimensional features and we suppose it will have a better performance than spatial segmentation on close pedestrians. Spectral Clustering is sensitive to the number of clusters (K in GMM), which needs to be set prudently. We directly initial the cluster number from ground truth to achieve the best performance. We also use the Munkres algorithm for data association.

For further discussion, we do not implement specific pedestrian tracking methods such as [2] [4] [5]. They discuss group people as an independent tracking element and do not try to separate them. These methods will be unfairly evaluated in our point-level annotation.

D. Evaluation: tracking

We evaluate the tracking performance of our method both qualitatively and quantitatively.

The qualitative results are demonstrated in Fig. 5. We analyze a typical group tracking example in our *hard* dataset.

The quantitative results are shown in Table I. We use several metrics [29] [30] which are widely used in multi-object tracking and can evaluate the precision and accuracy of a tracker. The selected MOT metrics are as follows:

- MOTP: The average errors between estimated positions and their corresponding ground truth.

TABLE I
EVALUATION ON TRACKING

Dataset	Method	MOTA \uparrow	MOTP \downarrow	FN \downarrow	FP \downarrow	IDS \downarrow
<i>easy</i>	EMST	0.800	0.038	213	33	3
	SC	0.936	0.083	25	29	37
	ours	0.999	0.004	0	2	0
<i>hard</i>	EMST	0.493	0.041	604	655	83
	SC	0.634	0.064	322	392	262
	ours	0.920	0.017	156	24	16

- MOTA: The ratio of correct data association and it indicates the ability to keep accurate trajectories.
- FP (false positives): The number of false alarms.
- FN (false negatives): The number of missed detections.
- IDS (identity switches): The number of mismatches.

Our method has the best performance in both datasets, especially in *hard* dataset which contains a certain amount of pedestrians in the scene and complex people interaction. We get a higher MOTA score of 0.920 in *hard* dataset, compared to the SC with 0.634 and the EMST with 0.493. In MOTP score, our method has an average error of 0.017m which is observably lower than the SC with 0.064m and the EMST with 0.041m. It should be noted that compared to FP and IDS, our method suffers more from FN. The *hard* dataset contains special instances that people crowd as a group in the whole sequence (Fig. 5). Our method sometimes clusters and tracks them as one object so the number of missed detections increases.

E. Evaluation: cardinality estimates

We also evaluate the cardinality estimates, which means the number of predicted pedestrians in each scan. The cardinality estimates of our method and the EMST baseline method are shown in Fig. 7.

Our error in pedestrian number is averagely 0.4557 per scan while the EMST has 0.4633. The difference is not significant between our method and the EMST baseline, but it is obvious in Fig. 7 that our tracking number changes more gently than the EMST method. This conclusion is also confirmed by the MOTA score in Table I.

F. Evaluation: timing

After ground and ROI filtering, the number of remained points is approximately 2000 which is stress-free for DB-SCAN. The number of iteration for EM optimization is averagely 3.6, which is much smaller than the maximum iterations we set of 10.

Our method is tested and evaluated on MATLAB R2019a. We also implement a ROS CPP version on a PC with an Intel i5-8600K CPU and 8 GB RAM. The consuming time per scan is averagely 51.47ms, which is sufficient for online use on our platform.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a pedestrian tracking framework in crowd scenarios based on Gaussian Mixture Model. The model is optimized by an extended Expectation-Maximization algorithm which can adaptively vary the number of mixture

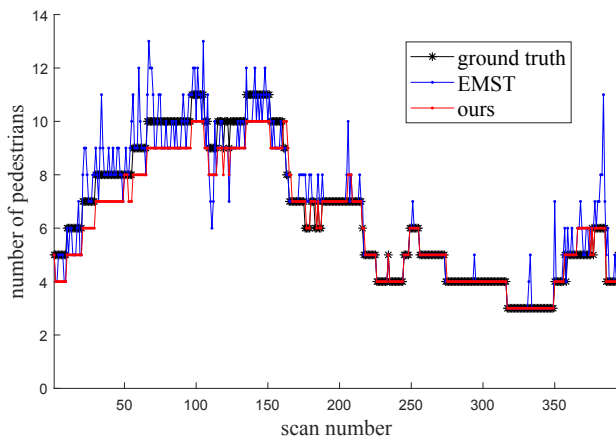


Fig. 7. Relationship between the cardinality estimates and the scan number for *hard* dataset sequence.

components over scans. Both qualitative and quantitative evaluation on tracking performance indicates that the proposed method shows robustness and accuracy in populated scenarios.

Our method shows good stability in pedestrian modeling. However, vehicles are sometimes contained in the scene which can be ambiguous in Gaussian representation. One extension of our work is to embed a 3D LiDAR detection CNN into our tracking module. Vehicles can be pre-treated and our framework can be expanded to a more general scenario.

REFERENCES

- [1] P. Zhang, W. Ouyang, P. Zhang, J. Xue, and N. Zheng, "Sr-lstm: State refinement for lstm towards pedestrian trajectory prediction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 12085–12094, 2019.
- [2] B. Lau, K. O. Arras, and W. Burgard, "Tracking groups of people with a multi-model hypothesis tracker," in *2009 IEEE International Conference on Robotics and Automation*, pp. 3180–3185, IEEE, 2009.
- [3] J. Shackleton, B. VanVoorst, and J. Hesch, "Tracking people with a 360-degree lidar," in *2010 7th IEEE International Conference on Advanced Video and Signal Based Surveillance*, pp. 420–426, IEEE, 2010.
- [4] G. Gate and F. Nashashibi, "Fast algorithm for pedestrian and group of pedestrians detection using a laser scanner," in *2009 IEEE Intelligent Vehicles Symposium*, pp. 1322–1327, IEEE, 2009.
- [5] M. Mucientes and W. Burgard, "Multiple hypothesis tracking of clusters of people," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 692–697, IEEE, 2006.
- [6] S. J. McKenna, S. Jabri, Z. Duric, A. Rosenfeld, and H. Wechsler, "Tracking groups of people," *Computer vision and image understanding*, vol. 80, no. 1, pp. 42–56, 2000.
- [7] A. Setia and A. Mittal, "Co-operative pedestrians group tracking in crowded scenes using an mst approach," in *2015 IEEE Winter Conference on Applications of Computer Vision*, pp. 102–108, IEEE, 2015.
- [8] R. Chandra, U. Bhattacharya, A. Bera, and D. Manocha, "Densepeds: Pedestrian tracking in dense crowds using front-rvo and sparse features," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 468–475, IEEE, 2019.
- [9] N. Bisagno, B. Zhang, and N. Conci, "Group lstm: Group trajectory prediction in crowded scenarios," in *Proceedings of the European conference on computer vision (ECCV)*, pp. 0–0, 2018.
- [10] D. Held, D. Guillory, B. Rebsamen, S. Thrun, and S. Savarese, "A probabilistic framework for real-time 3d segmentation using spatial, temporal, and semantic cues," in *Robotics: Science and Systems*, 2016.
- [11] J. Ku, A. D. Pon, S. Walsh, and S. L. Waslander, "Improving 3d object detection for pedestrians with virtual multi-view synthesis orientation estimation," *arXiv preprint arXiv:1907.06777*, 2019.

- [12] D. Z. Wang, I. Posner, and P. Newman, "What could move? finding cars, pedestrians and bicyclists in 3d laser data," in *2012 IEEE International Conference on Robotics and Automation*, pp. 4038–4044, IEEE, 2012.
- [13] K. Klasing, D. Wollherr, and M. Buss, "A clustering method for efficient segmentation of 3d laser data," in *2008 IEEE International Conference on Robotics and Automation*, pp. 4043–4048, IEEE, 2008.
- [14] I. Bogoslavskyi and C. Stachniss, "Fast range image-based segmentation of sparse 3d laser scans for online operation," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 163–169, IEEE, 2016.
- [15] A. Dewan, T. Caselitz, G. D. Tipaldi, and W. Burgard, "Motion-based detection and tracking in 3d lidar scans," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4508–4513, IEEE, 2016.
- [16] D. Z. Wang, I. Posner, and P. Newman, "Model-free detection and tracking of dynamic objects with 2d lidar," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 1039–1063, 2015.
- [17] A. Petrovskaya and S. Thrun, "Model based vehicle detection and tracking for autonomous urban driving," *Autonomous Robots*, vol. 26, no. 2-3, pp. 123–139, 2009.
- [18] B. Fortin, R. Lherbier, and J.-C. Noyer, "A model-based joint detection and tracking approach for multi-vehicle tracking with lidar sensor," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 4, pp. 1883–1895, 2015.
- [19] L. Spinello, K. O. Arras, R. Triebel, and R. Siegwart, "A layered approach to people detection in 3d range data," in *Twenty-fourth AAAI conference on artificial intelligence*, 2010.
- [20] M. He, E. Takeuchi, Y. Ninomiya, and S. Kato, "Precise and efficient model-based vehicle tracking method using rao-blackwellized and scaling series particle filters," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 117–124, IEEE, 2016.
- [22] R. Horaud, F. Forbes, M. Yguel, G. Dewaele, and J. Zhang, "Rigid and
- [21] H. Cho, Y.-W. Seo, B. V. Kumar, and R. R. Rajkumar, "A multi-sensor fusion system for moving object detection and tracking in urban driving environments," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1836–1843, IEEE, 2014.
- articulated point registration with expectation conditional maximization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 3, pp. 587–602, 2011.
- [23] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise.," in *Kdd*, vol. 96, pp. 226–231, 1996.
- [24] C. M. Bishop, *Pattern recognition and machine learning*. springer, 2006.
- [25] T. Chen, B. Dai, R. Wang, and D. Liu, "Gaussian-process-based real-time ground segmentation for autonomous land vehicles," *Journal of Intelligent and Robotic Systems*, vol. 76, no. 3, pp. 563–582, 2014.
- [26] J. Jiao, Y. Yu, Q. Liao, H. Ye, R. Fan, and M. Liu, "Automatic calibration of multiple 3d lidars in urban environments," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 15–20, Nov 2019.
- [27] J. Munkres, "Algorithms for the assignment and transportation problems," *Journal of the society for industrial and applied mathematics*, vol. 5, no. 1, pp. 32–38, 1957.
- [28] U. Von Luxburg, "A tutorial on spectral clustering," *Statistics and computing*, vol. 17, no. 4, pp. 395–416, 2007.
- [29] K. Bernardin, A. Elbs, and R. Stiefelwagen, "Multiple object tracking performance metrics and evaluation in a smart room environment," in *Sixth IEEE International Workshop on Visual Surveillance, in conjunction with ECCV*, vol. 90, p. 91, Citeseer, 2006.
- [30] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," *arXiv preprint arXiv:1903.11027*, 2019.